arXiv:1410.2568v1 [physics.data-an] 6 Oct 2014

Rising Above Chaotic Likelihoods

Hailiang Du[1,2] Leonard A. Smith[1,2,3]

[1]Centre for the Analysis of Time Series,

London School of Economics, London WC2A 2AE. UK

[2]Center for Robust Decision Making on Climate and Energy Policy,

University of Chicago, Chicago, IL 60637, US

[3]Pembroke College, Oxford, UK

October 10, 2014

Abstract

Berliner (Likelihood and Bayesian prediction for chaotic systems, J. Am. Stat. Assoc. 1991) identified a number of difficulties in using the likelihood function within the Bayesian paradigm for state estimation and parameter estimation of chaotic systems. Even when the equations of the system are given, he demonstrated "chaotic likelihood functions" of initial conditions and parameter values in the 1-D Logistic Map. Chaotic likelihood functions, while ultimately smooth, have such complicated small scale structure as to cast doubt on the possibility of identifying high likelihood estimates in practice. In this paper, the challenge of chaotic likelihoods is overcome by embedding the observations in a higher dimensional sequence-space, which is shown to allow good state estimation with finite computational power. An Importance Sampling approach is introduced, where Pseudo-orbit Data Assimilation is employed in the sequence-space in order first to identify relevant pseudo-orbits and then relevant trajectories. Estimates are identified with likelihoods orders of magnitude higher than those previously identified in the examples given by Berliner. Importance Sampling

uses the information from both system dynamics and observations. Using the relevant prior will, of course, eventually yield an accountable sample, but given the same computational resource this traditional approach would provide no high likelihood points at all. Berliner's central conclusion is supported. "chaotic likelihood functions" for parameter estimation still pose challenge; this fact is used to clarify why physical scientists tend to maintain a strong distinction between the initial condition uncertainty and parameter uncertainty.

1 Introduction

Nonlinear chaotic systems pose several challenges both for state estimation and for parameter estimation. Chaos as a phenomenon implies sensitive dependence on initial condition: initially nearby states will eventually diverge in the future. The bifurcations of various chaotic systems [20] reveal how the behavior of the system differs as a parameter value changes. One might think that likelihood and Bayesian analysis should be able to obtain good estimation both of initial conditions and of parameter values without much trouble. Berliner [2] examined the log-likelihood function of estimates of initial conditions and parameter values for the Logistic Map. He pointed out that chaotic systems can lead to "chaotic likelihood functions", suggesting that Bayesian analysis would require prohibitively intensive computing. The failure of variational approaches, when applied to long window observations of chaotic systems [6, 8, 14], supports his point. Sensitivity to initial condition also suggests that information in the observations (even over a relatively short range) can lead to good estimates of the initial condition [9]. An importance sampling approach for extracting such information without "intensive computing" is deployed in this paper. Adopting the Pseudo-orbit Data Assimilation (PDA) approach [6, 8] recasts the task into a higher dimensional sequence space, where truly high likelihood states are successfully located near the trajectory manifold. Although statisticians often fail to make a strong distinction between initial conditions and parameter values, the challenges of initial condition estimation and parameter estimation are dissimilar for chaotic systems. PDA

does not easily generalize to parameter estimation, as it is unclear how to mathematically define a relevant subspace of parameter space in which high likelihood trajectories might exist. Thus challenges remain in identifying high likelihood parameter values given the initial condition; this asymmetry is used to discuss differences between the initial conditions and parameter values. In terms of estimating initial conditions given the parameter values, however, Berliner's challenge is met and resolved without prohibitively intensive computing.

2 Chaotic Likelihood Function of Initial Conditions

Following Berliner [2] the Logistic Map is adopted as the system, assuming that the parameter $a = 4$ is known but the true initial state $\tilde{x}_0$ is not. In that case, the experiment is said to fall within the perfect model scenario[1].The evolution of system states $x_i \in \mathbb{R}^m$ is then governed by the nonlinear dynamics $f : x_{i+1} = f(x_i)$, where for the Logistic Map[2]

$$f(x_i) = ax_i(1 - x_i). \tag{1}$$

Assuming additive observational noise δ_i yields observations, $s_i = \tilde{x}_i + \delta_i$ where $\tilde{x}$ is the true system state (Truth) and the observational noise, δ_i, is Independent Normally Distributed (IND, $\delta_i \sim N(0, \sigma^2)$). Under this normality assumption, the log-likelihood (LLik) function is:

$$LLik(x_0) = -\sum_{i=1}^{n-1}(s_i - f^i(x_0))^2/2\sigma^2, \tag{2}$$

where f^i is the i^{th} iteration of f, s_i is the i^{th} observation, and n is the duration of observations considered.

Figure 1 shows the chaotic likelihood structure of 1024 samples from $U(0, 1)$. Panel (a) plots the log-likelihood for x_0, this can be contrasted with various panels in Berliner [2]

[1]By assuming a perfect model and generating the data on a digital computer, one avoid the issue of "round-off error": the fitted model is evaluated on the same computer. For discussion, see [1, 17, 18] and references thereof.

[2]For $m = 1$, the state x_i is a scalar.

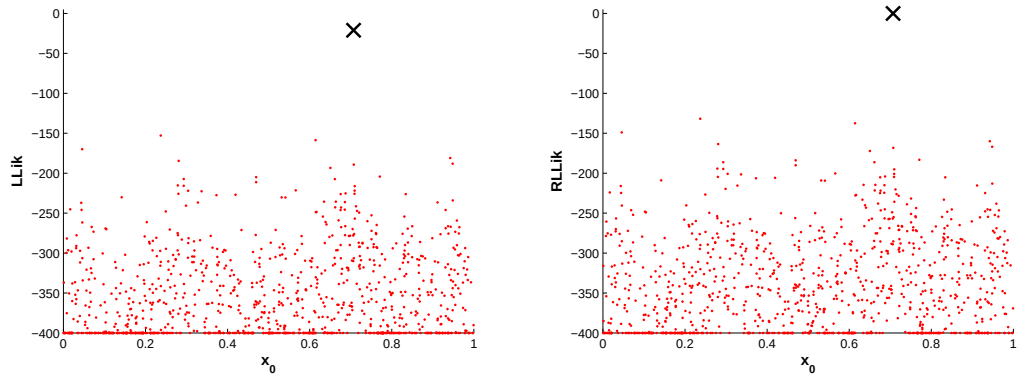

Figure 1: Typical log-likelihood of 1024 states (uniformly distributed on $[0, 1]$) for the Logistic Map. The true initial condition $\tilde{x}_0 = \sqrt{2}/2$, $\sigma = 0.1$ and $n = 32$. a) Log-likelihood function, b) Relative log-likelihood to $\tilde{x}_0$ (denoted by '$\times$'), states which have LLik/RLLik less than -400 are plotted on the -400 horizontal line. All logarithms are using natural base.

Figure 3[3]. Panel (b) shows the log-likelihood (RLLik) relative to that of the true trajectory of the system states[4]. For the convenience of illustration, the same normalization is applied in the following three figures in this paper. From Figure 1, it is clear that no high likelihood states are identified. This is not a case of equifinality[5].

Given the observational noise distribution, one can add random draws from the inverse of the observational noise distribution to the observation to obtain candidate estimates of initial condition. Figure 2a shows the relative log-likelihood of 1024 samples from inverse observational noise. No high likelihood states are identified in this way. To illustrate the

[3]Here LLik based on a sequence of 32 observations is computed. Berliner examined 15 & 10 observations. The problem becomes more obvious when more observations are used. Shorter sequences of observations are examined below.

[4]Note: given that only finite observations are considered the true state of the system is, with probability 1, not the maximum likelihood state

[5]Equifinality occurs when many potential solutions to a task are equally good, making it impossible to identify the true solution given the information in hand. In this case the sampled likelihood function is relatively flat. Equidismality arises when the sampled relative likelihood function is flat yet all solutions tested have vanishingly small likelihood given the information. Examining the relative likelihood obscures the difference; fortunately the expected (distribution of) likelihood can be computed from the noise model alone without knowledge of the true initial condition.

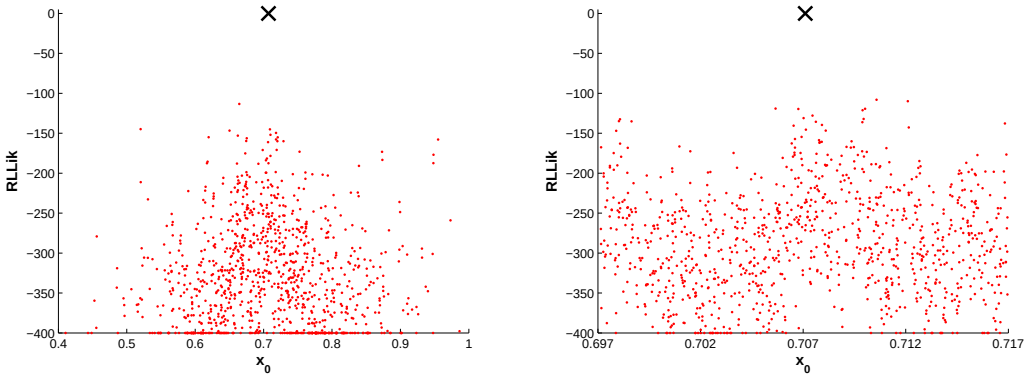

Figure 2: Log-likelihood of 1024 states for the Logistic Map, the true initial condition $\tilde{x}_0 = \sqrt{2}/2$, $\sigma = 0.1$ and $n = 32$. a) sampled from inverse observational noise, b) uniformly sampled from $[\tilde{x}_0 - \frac{\sigma}{10}, \tilde{x}_0 + \frac{\sigma}{10}]$

impact of making much more precise observations, consider a case where $\tilde{x}_0$ is known to be within a region of radius only $\sigma/10$. Figure 2b shows the RLLik of 1024 uniformly sampled states in the region around the Truth with $\sigma/10$ radius. Yet again, no high likelihood state are identified.

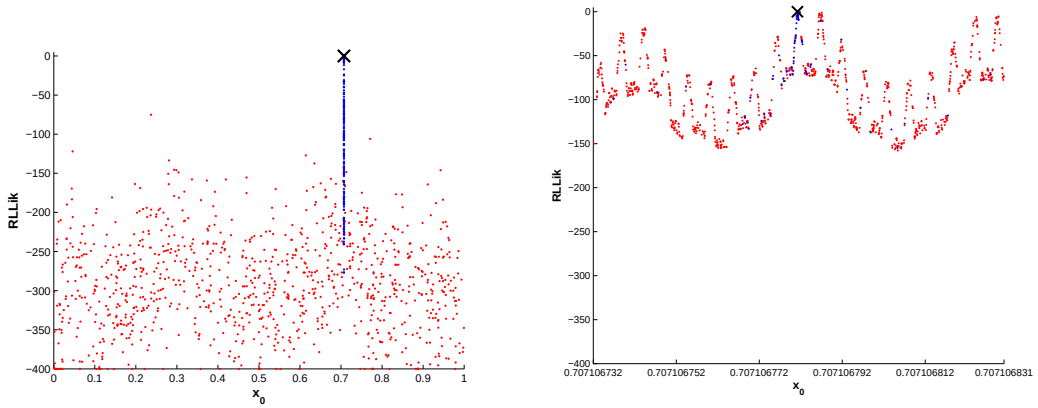

Figure 3: a) Following Figure 1b, add relative log-likelihood of 1024 states (blue), which are extremely close to $\tilde{x}_0$, generated by spiral sampling around the $\tilde{x}_0$; b) zoom in of a).

This difficulty here has nothing to do with the Likelihood approaches as there are high likelihood states other than Truth. One may demonstrate that such high likelihood states exist by sampling the points on a logarithmic spiral approaching the Truth (to machine

precision)[6]. Figure 3 shows that other than Truth there exist high likelihood states, i.e. some of the blue points. A smooth curve of the log-likelihood function is only observed within a radius of $\tilde{x}_0$ smaller than $\sim 10^{-7}$, see Figure 3b.

Without knowing the Truth, of course, this approach to identifying the blue points is inaccessible. The likelihood function is extremely jagged; as Berliner [2] stressed such chaotic likelihoods suggests that finding even one high likelihood state by sampling the state space would be prohibitively costly, making the approach inapplicable. That said there is no sense in which "sensitivity to the initial conditions" can be taken to imply that the information in the initial condition is "forgotten" or "lost". There is sufficient information in the observation segment to identify high likelihood initial states. Candidate states with non vanishing RLLik can be found by extracting the information from the system dynamics using a relatively new approach to data assimilation, which will be interpreted as an Importance Sampling.

3 Importance Sampling via Pseudo-orbit Data Assimilation

3.1 Methodology

To locate high likelihood states, simply sampling in state space is inefficient. As the dimension of the system increases, this inefficiency makes the task computationally impractical. Importance sampling[7] (IS) locates high likelihood states in the trajectory manifold by adopting the Pseudo-orbit Data Assimilation approach [6, 8]. PDA takes advantage of the known dynamics in a higher dimensional sequence space. A brief introduction of the PDA approach is given in the following paragraph (see [6, 8] for additional details).

[6]In this experiments 1024 points are generated by $\tilde{x}_0 + 2^{-(10+\frac{60i}{1024})}\epsilon_i, i = 1, 2, ..., 1024$ where ϵ_i is random drawn from $U(0, 1)$.

[7]In high-dimensional space the sampler targets the relevant low-dimensional trajectory manifold which is more efficient than sampling a hypersphere. Even in the one dimensional Logistic Map this approach succeeds by using PDA to sample the trajectory manifold in the n-dimensional sequence space.

Given a perfect dynamical model of dimension m, a perfect knowledge of the observational noise model, and a sequence of n observations $s_i, i = 0, ..., n-1$, define a sequence space as the $m \times n$ dimensional space consisting of all sequences of n states u_i.[8] Most points in sequence space do not correspond to a trajectory of the system. Define a *pseudo-orbit*, $\mathbf{U} \equiv \{u_0, ..., u_{n-2}, u_{n-1}\}$, to be a point in the $m \times n$ dimensional sequence space for which $u_{i+1} \neq f(u_i)$ for one or more components of $\mathbf{U}$. Thus a pseudo-orbit corresponds to a sequence of system states which is not a trajectory of the system. Define the *mismatch* to be:

$$e_i = \mid f(u_i) - u_{i+1} \mid \tag{3}$$

By construction, system trajectories have a mismatch of zero. The mismatch cost function is then given by:

$$C(\mathbf{U}) = \sum_{i=0}^{n-1} e_i^2 \tag{4}$$

Pseudo-orbit Data Assimilation minimizes the mismatch cost function for $\mathbf{U}$ in the $m \times n$ dimensional sequence space. If a gradient descent (GD) approach is adopted[9], then a minimum of the mismatch cost function can be obtained by solving the ordinary differential equation

$$\frac{d\mathbf{U}}{d\tau} = -\nabla C(\mathbf{U}), \tag{5}$$

where τ denotes algorithmic time.[10] A sequence of observations in the system state space define an initial pseudo-orbit, so called *observation-based pseudo-orbit*, $\mathbf{S} \equiv \{s_0, ..., s_{n-2}, s_{n-1}\}$, which with probability one will not be a trajectory. In practice, the minimization is initialized with the observation-based pseudo-orbit, i.e. ${}^0\mathbf{U} = \mathbf{S}$ where the pre-super-script ^{0} on $\mathbf{U}$ denotes the initial stage of the GD. The pseudo-orbit is a point in sequence space,

[8] For the Logistic Map where $m = 1$, both s_i and u_i are scalars.

[9] Other methods for this minimization are available, GD is discussed here due to its simplicity.

[10] The approach can be generalized to situations where gradient of the model is not known analytically [10]; improving the ability to work without gradient information would widen the applications of the approach significantly.

under Equation 5 this point moves towards the manifold of all trajectories. The mismatch cost function has no local minima other than on the manifold, for which $C(\mathbf{U}) = 0$, (i.e. the trajectory manifold[11]) and every segment of trajectory lies on this manifold [8]. Let the result of the GD minimization at time α be ${}^{\alpha}\mathbf{U}$. Here α indicates algorithmic time in GD (i.e. the number of iterations of the GD minimization). As $\alpha \to \infty$, the pseudo-orbit ${}^{\alpha}\mathbf{U} \equiv {}^{\alpha}\ \mathbf{u}_0, ..., {}^{\alpha}\ \mathbf{u}_{n-1}$ approaches a trajectory of the model asymptotically. In other words, the GD minimization takes us from the observation-based pseudo-orbit towards a system trajectory (a point in sequence space, ${}^{\infty}\mathbf{U}$, which is on the trajectory manifold). In practice, the GD minimization is run for a finite time and thus a trajectory is not obtained. The result of these large α GD runs, ${}^{\alpha}u_0$ provide candidates for the initial state, based on information from the observations with $i < n$. For $i > 0$, the i-step preimage of the relevant component of ${}^{\alpha}u_i$ are calculated to obtain additional candidates for the initial state. The Logistic Map is a two-to-one map, and in most cases[12] only one of the two preimages for each ${}^{\alpha}u_i$ is relevant to $\tilde{x}_{i-1}$. In practice a criteria to discard irrelevant preimages must be defined, a simple example would be to discard (with high probability) those preimages whose distance from the corresponding previous observation exceeds some threshold based on the standard deviation of the observational noise (a 3σ criteria is used to generate the results presented in the following section).

3.2 Results

The green points in Figure 4 are located using the IS approach; Note that some have RLLik close to 0. As expected, the observations do not contain sufficient information to identify the state of the system at the time of the final observation with the same degree of precision. This is reflected in the fact that the green points are much less close to the true state at time 31 (Figure 4b) than those at time 0 (Figure 4a).

Two experiments were conducted to test the robustness of the IS approach. The first is

[11] All points on the trajectory manifold have zero mismatch (are trajectories) and only points on the trajectory manifold have zero mismatch.

[12] Not in all cases, however. For discussion of the point see [11].

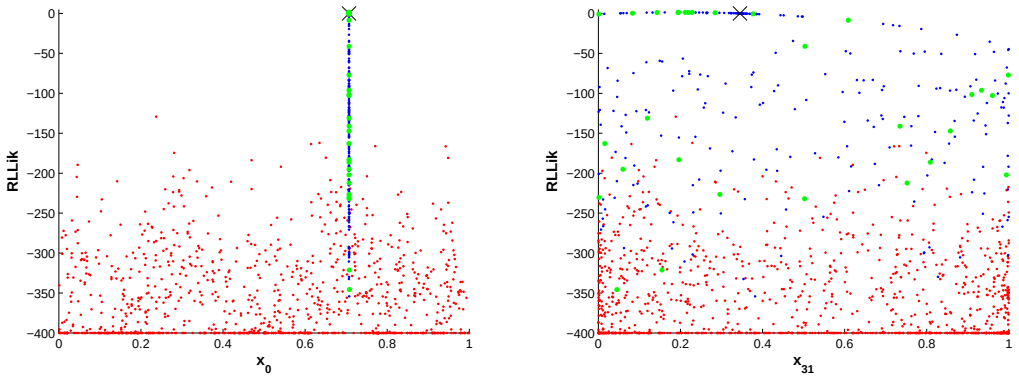

Figure 4: Following Figure 3a, Relative log-likelihood of the states located by IS are plotted in green, a) $i = 0$; b) $i = 31$.

based on 2048 different realizations of observations for $\tilde{x}_0 = \sqrt{2}/2$ to examine consistency. The second is based on 2048 different true initial conditions to examine robustness. Three different observation window lengths were used in each experiment. Table 1 and Table 2 shows the results.

Given uncertain observations, one can never identify the Truth of a chaotic system unambiguously as was noted by Lalley [11, 12] and later explored by Judd and Smith [8]. Using the IS approach, high likelihood states (IS states) are indeed found, as the states whose RLLik> -1 are found in every single experimental run. The fact that some IS states have greater likelihood than the Truth supports the expectation that the Truth is not expected to be the most likely system state given the observations.

For each experimental run, the minimum distance between those IS states (whose RLLik> -1) and the Truth is recorded. The minimum, maximum and median statistical values of the minimum distance from the Truth are reported in Table 1 and 2. It is clear that the quality of the IS states improves (the minimum distance from the Truth decreases) as the observation window length increases. This is expected inasmuch as more information from the system dynamics becomes available when using a longer window. It is shown in Table 1 that the maximum value of the minimum distance among the 2048 different realizations is 1.49×10^{-10} for a window length of 32 and in Table 2 the maximum value of the minimum distance among different true initial conditions is 2.02×10^{-10}. PDA

importance sampling appears both robust and efficient[13].

Window	# of RLLik> -1			# of RLLik> 0			Minimum distance to $\tilde{x}_0$		
length	Min	Max	Median	Min	Max	Median	Min	Max	Median
32	6	15	8	0	14	6	2.00×10^{-15}	1.49×10^{-10}	9.98×10^{-12}
16	2	11	8	0	11	6	1.57×10^{-10}	7.63×10^{-6}	5.13×10^{-7}
8	2	7	7	0	7	5	8.58×10^{-8}	4.55×10^{-2}	2.56×10^{-4}

Table 1: Statistics of high likelihood states located by IS based on 2048 different realizations of observations (of $\tilde{x}_0 = \sqrt{2}/2$) for the Logistic Map, i) statistics of the number of states (whose $RLLik > -1$) ii) statistics of the number of states (whose $RLLik > 0$) iii) statistics of the minimum distance between the states (whose $RLLik > -1$) and the Truth.

Window	# of RLLik> -1			# of RLLik> 0			Minimum distance to $\tilde{x}_0$		
length	Min	Max	Median	Min	Max	Median	Min	Max	Median
32	6	360	28	0	338	16	4.77×10^{-15}	2.02×10^{-10}	1.50×10^{-11}
16	3	56	14	0	48	9	3.04×10^{-10}	1.54×10^{-5}	9.93×10^{-7}
8	2	20	7	0	20	7	6.36×10^{-8}	5.60×10^{-3}	3.32×10^{-4}

Table 2: Statistics of high likelihood states located by IS based on 2048 different true initial states for the Logistic Map, i) statistics of the number of states (whose $RLLik > -1$) ii) statistics of the number of states (whose $RLLik > 0$) iii) statistics of the minimum distance between the states (whose $RLLik > -1$) and the Truth.

Complications arise from the fact that the Logistic Map is two-to-one; these have nothing to do with chaos per se (beyond the fact that one-to-one maps in one-dimension cannot display chaotic dynamics). Moving to higher dimensional[14] one-to-one maps, the calculation of preimages becomes straightforward.

[13]Drawing samples uniformly from within a distance of 0.1 of Truth would require $\sim 10^8$ candidates in order to find a candidate within $\sim 2 \times 10^{-10}$ of Truth. Such a procedure need not identify any high likelihood states. The results of Table 1 and 2 were obtained with only 1024 GD minimization iterations in each realization (each and every one of which identified high likelihood states close to Truth).

[14]where the model state becomes a state vector

The experiments above demonstrate that truly high likelihood points can be located using dynamical information. This eases Berliner's identification problem of initial condition with the appearance of chaotic likelihoods. Selecting an ensemble from this high likelihood set allows for informative forecasts which do not become useless until long after those from the point forecasts illustrated by Berliner [2] become uninformative.

4 Relative likelihoods

Maximum Likelihood Estimation has been widely used for estimation [15] since it was introduced by Fisher [7] in 1922. The "best" estimate is often chosen from a set of samples and only the relative likelihood in that sample is considered. Figure 5a shows the log-likelihood of 1024 states (the same set used in Figure 2b), the grey dashed line is the median log-likelihood of those states. In this case, it is not the problem of which estimate one shall pick, but how to show that they are all "bad" estimates. In practice, the Truth is unknown therefore it cannot be used as a reference like the cross plotted in Figure 2. Given the observations and the noise model, however, the expected log-likelihood of the Truth[15] can be derived and serve as a reference. Figure 5b, the log-likelihood of 1024 states are plotted along with the expected log-likelihood of the Truth (black dashed line). Figure 5b shows that it is not a case of equifinality in Figure 5a but a case of equidismality. In cases where it is observed that all traditional candidate states have vanishingly small log-likelihood relative to the expected log-likelihood of the Truth, approaches like those suggested above might prove valuable.

[15]The log-likelihood of the Truth is $\frac{\sum_{i=0}^{n-1}\delta_i^2}{2\sigma^2}$ (from Eq. 2) where δ_i (observation noise) is $IID \sim N(0,\sigma^2)$ distributed. Let $Z = \frac{\sum_{i=0}^{n-1}\delta_i^2}{\sigma^2}$, Z is a random variable following chi-squared distribution with n degrees of freedom. Statistics of the log-likelihood of the Truth can therefore simply derived from Z.

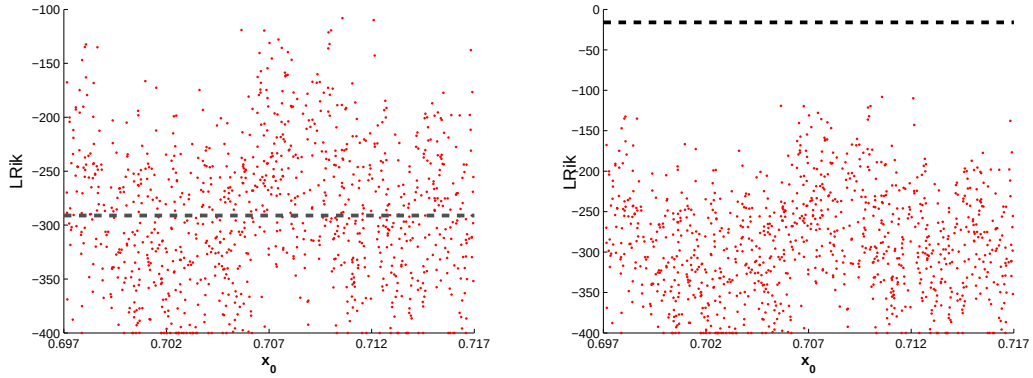

Figure 5: Following Figure 2b Log-likelihoods of 1024 states uniformly sampled from $[\tilde{x}_0 - \frac{\sigma}{10}, \tilde{x}_0 + \frac{\sigma}{10}]$, a) the grey dashed line is the median log-likelihood of those states, b) the black dashed line is the expected log-likelihood of the Truth

5 Difference between initial conditions and parameters

Statisticians often treat estimating initial conditions and estimating parameter values as similar problems. Although similar behaviors of likelihood functions of initial conditions and that of parameter values are observed [2], there are fundamental differences in the information available to address these two distinct estimation problems.

Given the structure of the model class, the model parameter value determines the dynamical behaviour of the model (e.g. the natural measure) which is not changed by the initial condition. Given the model and its parameter value(s), the invariant measure constrains the relevant set(s) of initial conditions in the state space (and thereby trajectories in the sequence space). It is unclear how to construct similar constraints (if they exist[16]) on the parameter values in the parameter space given the "true" initial state. Uncertainty in initial state differs from uncertainty in the parameter value. The information in a measurement of the initial condition uncertainty will decay with time and eventually becomes

[16]It is not clear either how to construct the set of parameter values whose corresponding invariant measure contains the "true" initial state, or how to exploit this set, while it is clear how to exploit the existence of trajectory manifold given a particular value of the parameter.

statistically indistinguishable from a random sample of the natural measure, while the information on each member from an ensemble under parameter uncertainty is preserved (and can be straight forwardly extracted given a trajectory segment), arguably forever.

While assuming the parameter value is perfect may not be ideal, it is not so nonsensical given that one has already assumed that the model structure is perfect. Assuming the initial state is perfect indicates a noise free observation is possible. Let the model's parameters be contained in the vector $\mathbf{a} \in \mathbb{R}^l$. A set of $l+1$ sequential noise free observations $s_i, s_{i+1}, ..., s_{i+l}$ would, in general, be sufficient to determine $\mathbf{a}$ [13]. If one noise free observation is obtainable, obtaining only a few more noise free observations would define the true parameter value precisely. A more realistic way to put the problem is to estimate the parameter value(s) given the observations without assuming the "true" initial condition is known or even exists. In that case, the goal is to locate high likelihood trajectories (Smith et al. [19] call these shadowing trajectories) defined by the parameter values. Unfortunately it is not clear how to solve such a problem. In fact, it is not clear how to constrain the solution in the parameter space in a manner that reflects the constraints in the space of initial condition achieved by using trajectory manifold in the state space.

Given a perfect model structure and knowing the true parameter value(s), the true initial state is a well defined goal of the identification. Inasmuch as structural model errors imply no true parameter value exists [5, 9], it is unclear how one might define "true" initial state and the goal of estimation must be rethought.

Despite the importance of model parameters, there is no general method of parameter estimation outside linear systems. Methods have been developed to obtain useful parameter values with some success: McSharry and Smith [13, 16] estimate model parameters by incorporating the global behaviour of the model into the selection criteria; Creveling et al. [3] have exploited synchronization for parameter estimation; Smith et al. [19] focused on the geometric properties of trajectories; Du and Smith [4] select parameter values based on the Ignorance Score of ensemble forecasts. Each of these methods, however, require a large set of observations. Challenges remain when only a short sequence of observations is available.

6 Conclusion

Berliner illustrated that even in the perfect model scenario traditional likelihood methods are unable to provide good estimates of the initial condition for nonlinear chaotic systems. In large part, the failure is due to the inability of those approaches to skillfully meld the information in the dynamics of the nonlinear system itself with that in the observations. The importance sampling approach presented here better combines information from both observations and dynamics, thereby locating high likelihood initial states; this achieves an aim Berliner (1991) argued to be impossible by traditional methods. Despite the similarity of state estimation and parameter estimation, there are fundamental differences between uncertainty in the initial state and uncertainty in parameter value. Significant challenges remain in solving the challenge chaotic likelihood functions pose in parameter estimation.

Acknowledgment

This research was supported by the LSE's Grantham Research Institute on Climate Change and the Environment and the ESRC Centre for Climate Change Economics and Policy, funded by the Economic and Social Research Council and Munich Re. Additional support for H.D. was also provided by the National Science Foundation Award No. 0951576 "DMUU: Center for Robust Decision Making on Climate and Energy Policy (RDCEP)". L.A. S. gratefully acknowledges the continuing support of Pembroke College, Oxford.

References

[1] C. Beck. Effects of roundoff errors on chaotic dynamics. In J. Vandewalle, R. Boite, M. Moonen, and A. Oosterlinck, editors, Signal Processing VI - Theories and Applications Proceedings of EUSIPCO-92, volume 1, page 183186, 1992.

[2] M.L. Berliner. Likelihood and bayesian prediction for chaotic systems. J. Am. Stat. Assoc., 86:938952, 1991.

[3] D.R. Creveling, P.E. Gill, and H.D.I. Abarbanel. State and parameter 258 estimation in nonlinear systems as an optimal tracking problem. Physics Letters A, 372:26402644, 2008.

[4] H. Du and L.A. Smith. Parameter estimation through ignorance. Phys. Rev. E, 016213, 2012.

[5] H. Du and L.A. Smith. Pseudo-orbit data assimilation. II: Assimilation with imperfect models. Journal of the Atmospheric Sciences, 71:483495, 2014.

[6] H. Du and L.A. Smith. Pseudo-orbit data assimilation part I: The perfect model scenario. Journal of the Atmospheric Sciences, 71:469482, 2014.

[7] R.A. Fisher. On the mathematical foundations of theoretical statistics. Philosophical Transactions of the Royal Society of London, 222:309368, 1922.

[8] K. Judd and L.A. Smith. Indistinguishable states I: The perfect model scenario. Physica D, 151:125141, 2001.

[9] K. Judd and L.A. Smith. Indistinguishable states II: the imperfect model scenario. Physica D, 196:224242, 2004.

[10] Kevin Judd, Leonard Smith, and Antje Weisheimer. Gradient free descent: shadowing, and state estimation using limited derivative information. Physica D: Nonlinear Phenomena, 190(34):153166, April 2004.

[11] S.P. Lalley. Beneath the noise, chaos. The Annals of Statistics, 27(2):461479, 1999.

[12] S.P. Lalley. Removing the noise from chaos plus noise. In A.I. Mees, editor, Nonlinear Dynamics and Statistics, Boston, 2000.

[13] P. McSharry and L.A. Smith. Better nonlinear models from noisy data: Attractors with maximum likelihood. Phys. Rev. Lett, 83(21):42854288, 1999.

[14] R.N. Miller, M. Ghil, and F. Gauthiez. Advanced data assimilation in strongly nonlinear dynamical systems. J. Atmos. Sci., 51:10371056, 1994.

[15] Y. Pawitan. In all likelihood : statistical modelling and inference 283 using likelihood. Oxford science publications. Clarendon press, Oxford, 2001.

[16] V.F. Pisarenko and D. Sornette. Statistical methods of parameter estimation for deterministically chaotic time series. Phys. Rev. E, 69:036122, 2004.

[17] L.A. Smith. Lacunarity and Chaos in Nature. PhD thesis, Columbia University, New York, 1987.

[18] L.A. Smith. What might we learn from climate forecasts? In Proc. National Acad. Sci., volume 4, pages 2487-2492, USA, 2002.

[19] L.A. Smith, M.C. Cuellar, H. Du, and K. Judd. Exploiting dynamical coherence: A geometric approach to parameter estimation in nonlinear models. Physics Letters A, 374:26182623, 2010.

[20] J.C. Sprott. Chaos and Time-Series Analysis. Oxford University Press, Oxford, 2003.